

"Express Mail" mailing label number:

EL764879881US

A METHOD FOR RESTORING A VIRTUAL PATH IN AN OPTICAL NETWORK USING DYNAMIC UNICAST

Saleh, Ali Najib; Zadikian, H. Michael; Baghdasarian, Zareh; Parsi, Vahid

5

CROSS-REFERENCES TO RELATED APPLICATIONS

This application is a continuation-in-part of Patent Application Serial No. 09/858,743, filed May 16, 2001 and entitled "A Resource Reservation Scheme For Path Restoration In An Optical Network," having A. N. Saleh, H. M. Zadikian, Z. Baghdasarian, and V. Parsi as inventors. This application is assigned to Cisco Technology, Inc the assignee of the present invention, and is hereby incorporated by reference, in its entirety and for all purposes.

BACKGROUND OF THE INVENTION

Field of the Invention

This invention relates to the field of information networks, and more particularly relates to a protocol for configuring routes over a network.

Description of the Related Art

Today's networks carry vast amounts of information. High bandwidth applications supported by these networks include streaming video, streaming audio, and large aggregations of voice traffic. In the future, these bandwidth demands are certain to increase. To meet such demands, an increasingly popular alternative is the use of lightwave communications carried over fiber-optic cables. The use of lightwave communications provides several benefits, including high bandwidth, ease of installation, and capacity for future growth.

Optical infrastructures are capable of transmission speeds in the gigabit range, which helps address the ever-increasing need for bandwidth mentioned above. Such infrastructures employ various topologies, including ring and mesh topologies. In order to provide fault protection, ring topologies normally reserve a large portion (e.g. 5 50% or more) of the network's available bandwidth for use in restoring failed circuits. However, ring topologies are capable of quickly restoring failed circuits. This capability is important in providing reliable service to customers, and is particularly important in telephony applications, where a failure can result in alarms, dropped calls, and, ultimately, customer dissatisfaction and lost revenue. In a similar vein, 10 because of bandwidth demands, protocol overhead related to provisioning, restoration, and other functions should be kept to a minimum in order to make the maximum amount of bandwidth available for use by customers.

An alternative to the ring topology, the mesh topology reduces the amount of bandwidth needed for protection. The mesh topology is a point-to-point topology, 15 with each node in the network connected to one or more other nodes. Because a circuit may be routed through various combinations of the network's nodes and over the various links which connect them, excess capacity through a given node or over a given link can serve to protect several circuits. However, the restoration of a circuit following a failure in a mesh topology can consume a relatively large amount of time.

20

SUMMARY

In one embodiment of the present invention, a method is described for restoring a virtual path in an optical network. The method includes identifying one or 25 more nodes with resources, wherein the nodes with resources are ones of the nodes having a resource necessary to support the virtual path, determining an alternate physical path, the alternate physical path includes the nodes with resources, and restoring the virtual path using the alternate physical path. The restoring is done by configuring the alternate physical path by establishing a communication connection 30 between the nodes with resources and provisioning the virtual path over the alternate physical path.

In another embodiment, the method includes detecting a failure in the virtual path wherein the failure is detected by receiving a failure message packet.

The foregoing is a summary and thus contains, by necessity, simplifications, generalizations and omissions of detail; consequently, those skilled in the art will appreciate that the summary is illustrative only and is not intended to be in any way 5 limiting. Other aspects, inventive features, and advantages of the present invention, as defined solely by the claims, will become apparent in the non-limiting detailed description set forth below.

BRIEF DESCRIPTION OF THE DRAWINGS

10

The present invention may be better understood, and numerous objects, features, and advantages made apparent to those skilled in the art by referencing the accompanying drawing.

Fig. 1 is a block diagram of a zoned network.

15

Fig. 2 illustrates an example of Add Path Request packet.

Fig. 3 illustrates a flow chart of steps performed by a tandem node when processing an Add Path Request.

Fig. 4 illustrates a flow chart of steps performed by a destination node when processing an Add Path Request.

20

Fig. 5 illustrates an example of Add Path Response packet.

Fig. 6A is a flow chart illustrating the actions performed by a source node when the source node receives a response packet.

Fig. 6B is a flow chart illustrating the actions performed by a tandem node when the tandem node receives a response packet.

25

Fig. 7 illustrates an example of a virtual path restoration in a network

DETAILED DESCRIPTION OF THE INVENTION

The following is intended to provide a detailed description of an example of 30 the invention and should not be taken to be limiting of the invention itself. Rather, any number of variations may fall within the scope of the invention which is defined in the claims following the description.

Introduction

A network can employ various restoration schemes to restore a virtual path
5 (VP) in case of a failure. To guarantee the restoration of a VP in, each VP is assigned
a restoration priority level. The restoration priority level determines a VP's relative
priority with regard to restoration in the event of a failure within the network. The
present invention provides a method of restoring a virtual path using dynamic unicast.
In such a dynamic unicast method, a VP is assigned a single path when provisioned.
10 The restoration of the VP is guaranteed. The VP is restored by creating a new physical
path and provisioning the virtual path on the new physical path.

Network Configuration

15 Fig. 1 illustrates an exemplary zoned mesh network. The network can be
configured in a non-zoned network in which network nodes are coupled in a mesh
topology. The exemplary zoned mesh network of Fig. 1 has been organized into a
backbone, zone 100, and four configured zones, zones 101-104. The solid circles in
each zone represent network nodes, while the numbers within the circles represent
20 node addresses, and include network nodes 111-117, 121-126, 131-136, and 141-147.
The dashed circles represent network zones. The network depicted in Fig. 1 has four
configured zones (zones 101-104 (addressed as zone 1-4) and one backbone (zone
100). Network nodes 113, 117, 122, 124, 134, 135, 141, and 142, are boundary or
proxy nodes because they connect to more than one zone. All other nodes are interior
25 nodes because their links attach only to nodes within the same zone. However, the
exemplary network of Fig. 1 can be configured as non-zoned mesh network. In a non-
zoned mesh network, nodes are combined into one network with no boundary or
proxy node.

Provisioning of network nodes

30 Once a mesh network topology has been defined (e.g., the zoned topology of
Fig. 1), the user can configure one or more end-to-end connections that can span
multiple nodes or zones, an operation is referred to herein as provisioning. For each

virtual path to be provisioned, a physical path must be selected and configured. Each set of physical connections that are provisioned creates an end-to-end connection between the two end nodes that supports a virtual point-to-point link (referred to herein as a virtual path or VP). The resulting VP has an associated capacity and an

5 operational state, among other attributes.

In a network, VPs may be provisioned statically or dynamically. For example, a user can identify the nodes that will comprise the virtual path and manually configure each node to support the given virtual path. The selection of nodes may be based on any number of criteria, such as Quality of Service (QoS), latency, cost, distance traveled in the network and the like. Alternatively, the VP may be provisioned dynamically using any one of a number of methods. The provisioning information may then be forwarded to all the nodes in the network to store information in node's network topology database. Each node periodically updates this information to efficiently maintain resources and in case of path failure, effectively allocate appropriate resources needed for specific virtual path for path restoration. The method of routing information in such networks is described in a commonly-assigned U.S. Patent Application Serial No. 09/232,395, entitled "A Configurable Network Router," filed 1/15/99, which is hereby incorporated by reference, in its entirety and for all purposes.

20 The end nodes of a VP can be configured to have a master/slave relationship. The terms source and destination are also used herein in referring to the two end-nodes. In such a relationship, the node with a numerically lower node ID typically assumes the role of the master (or source) node, while the other assumes the role of the slave (or destination) node, although the opposite arrangement is also acceptable.

25 An intermediate node is referred to herein as tandem node. Typically, the source node assumes the provisioning responsibilities and the destination node simply waits for a message from the source node informing the destination node of the VP's new physical path (although again, this need not necessarily be the case). This information includes node identifiers of tandem nodes, if any, within the path. In a zoned mesh topology, if a virtual path spans over multiple zones, the border node or proxy node of each zone acts as source node for their particular zone. As will be apparent to one of skill in the art, the opposite convention or another paradigm can easily be employed.

Typically, during provisioning, each VP is assigned a performance and restoration priority level. The priority, referred to herein as Class of Service (CoS), determines VP's relative priority for performance within the network and restoration in the event of a failure within the network. The method of assigning CoS to a VP is 5 described in commonly-assigned U.S. Patent Application Serial No. 09/858,743, filed on May 16, 2001, entitled "A Resource Reservation Scheme for Path Restoration in an Optical Network," which is hereby incorporated by reference, in its entirety and for all purposes. In case of a VP failure at a node in the network, the node determines how to restore the VP based on the CoS assigned to the VP. The assigned CoS defines the restoration method used by the node to restore failed VP.

10

Failure detection, propagation, and path restoration

Failure Detection and Propagation

15

In networks, failures are typically detected using the mechanisms provided by the underlying physical network. The failure detection mechanism in a mesh optical network is described in commonly-assigned U.S. Patent Application, Serial No. 09/232,397, filed 1/15/99 and entitled "A Method For Routing Information Over A 20 Network," which is hereby incorporated by reference, in its entirety and for all purposes.

Path Restoration using Dynamic Unicast Method

If the CoS of a VP defines dynamic unicast as the method of restoration for that VP, then the restoration is guaranteed but the restoration time can be longer than 25 what might be needed for critical path traffic such as a voice call. The "restoration time guarantees" for a VP using dynamic unicast method depend on the network configuration and provisioning attributes (e.g., bandwidth requirement, QoS, latency and the like) of the VP. In a zoned network topology, the dynamic unicast restoration method is generally used to restore intra-zone path failures. Preferably, the source 30 node of failed VP initiates the dynamic unicast restoration process. The tandem node that discovers the path failure, initiates a path failure notification for the source node and waits for a response. The destination node preferably responds to restoration process initiated by the source node.

Initiating restoration

Once a node other than the source node detects a path failure, the node initiates a path restoration request for the source node of failed VP using a *Restore_I* request. The method of generating *Restore_I* requests and responses is described in 5 commonly-assigned U.S. Patent Application Serial No. 09/750,668, entitled “A Virtual Path Restoration Scheme Using Fast Dynamic Mesh Restoration in an Optical Network”, filed on 12/29/00 and is hereby incorporated by reference in its entirety and for all purposes.

When the source node receives a *Restore_I* request, the source node 10 determines the type of restoration scheme assigned to the failed VP. Upon determining that the failed VP is to use the dynamic unicast restoration scheme, the source node creates an Add Path Request packet with appropriate contents and transmits the request to other nodes in the network.

Add Path Request Packet

Fig. 2 illustrates an example of an Add Path Request (APR) packet 200. APR 15 packet 200 is also used to provision a new path for a VP. APR packet 200 includes a VP-ID 210, the 32 bit ID of the VP. It will be apparent to one of the skill in the art that, while specific lengths are described, the fields discussed here may be of any appropriate length. A 2-bit long ‘T’ field 220 is used to indicate the type of path. 20 This field indicates whether the path is primary path or secondary path for a CoS 3 VP. A 2-bit long ‘C’ field 230 defines the restoration Class of Service. A 12-bit long bandwidth field 240 indicates the bandwidth requested for the VP for example, in STS-48 granularity.

A physical instance field 250 stores a 16-bit physical instance identifier for the 25 VP. The source node of the VP maintains physical instance field 250, which is associated with the path (i.e., set of link IDs) of the VP and is part of APR packet 200 and all restoration-related packets. Preferably, the source node updates physical instance field 250. The first path of a VP that is successfully provisioned (as seen by the source node) has a physical instance identifier of 1. All future path messages 30 should have the correct value of this identifier. If a new path is selected for the VP, the physical instance identifier is incremented (e.g., by 1) by the source node. Due to

the distributed nature of path selection and multiple failures, several physical instances of the same VP may temporarily exist in the network at any given time. However, only one instance ultimately survives.

An attempt count field 260 is the attempt count of the current physical

- 5 instance of the VP. The source node increments this field every time the source node resends the same request. Since APR packets are retransmitted periodically, different attempts preferably should be distinguished from one another. Attempt count field 260 allows retransmitted requests to be distinguished from one another. Attempt count field 260 starts at a given point (e.g., from 1) and is incremented (e.g., by 1)
10 with each retransmission. Because given APR packets may traverse different paths to get to the same intermediate node, attempt count field 260 allows the intermediate node to differentiate among multiple request attempts for restoration of the same physical instance of the same VP.

A path length field 270 indicates the number of links in the VP. Path length

- 15 field 270 determines the number of link identifications that appear at the end of the packet. A hops field 280 indicates the number of hops traversed by a given APR packet. Hops field 280 is incremented (e.g., by 1) at each receiving node in the given APR packet. During the transmission of the given APR packet, the value of hops field 280 is incremented (e.g., from 0 to (Path Length - 1)). Upon the return of a
20 response, hops field 280 is decremented (e.g., by 1) by each node that forwards the response to the source node. During the transmission of a response, the value of hops field 280 is decremented from the maximum number of hops traversed to zero by the time the response reaches the source node (e.g., from (Path Length - 1) to 0). A link ID field 290 is a 32-bit long field for the link IDs of the VP. The number of link IDs
25 depends upon the path length set by the source node.

Upon sending APR packet 200, the source node sets a timer. If a positive response is not received before the timer expires, the source node generates another APR packet. Each time the APR packet is generated, attempt count field 260 is incremented (e.g., by 1). The APR packets are preferably generated for only a certain number of times, after which the source node generates a network alarm.
30

Receiving an Add Path Request packet at a Tandem Node

Fig. 3 is a flow chart illustrating the actions performed by a tandem node upon receipt of an APR packet. Initially, the tandem node receives an APR packet (step 305). The tandem node first determines if the request is invalid and contains any error (step 310). If the request is valid, the node proceeds with resource checking (step 360). Otherwise, the process makes at least one of several determinations as to the reason for the request's invalidity. The tandem node determines if the APR packet is misrouted (i.e., there is no path to the destination, or the tandem node does not recognize any of the link IDs in the list due to an incomplete network topology update) (step 315). If the APR packet is misrouted, the tandem node returns a NAK (UNREACHABLE) upstream (step 320). If the APR packet is not misrouted, the tandem node determines if the path type field in the APR packet is invalid (step 325). If the path type field is invalid, the tandem node responds with a NAK (INVALID PATH TYPE) (step 330). If the path type field is valid, the tandem node determines if the APR packet is an unexpected request (i.e., the request arrives while the path is in restoring or deleting states) (step 340). If the APR packet is an unexpected request, the tandem node responds with a NAK (WRONG STATE) (step 345). If the APR packet is not unexpected, the tandem node determines if one or more of the parameters of the APR packet (such as the CoS, origin, target or bandwidth fields) are invalid (step 350). If one or more of the parameters of the APR packet are invalid, the tandem node returns a NAK (NO RESOURCES) to the source node (step 355).

If the APR packet is valid and no errors are found, the tandem node determines if sufficient resources are available to support the virtual path (step 360). If there are insufficient resources (such as memory, bandwidth on input and/or output links, unavailability of ports and the like), the tandem node responds with a NAK (NO RESOURCES) (step 355). However, if sufficient resources are available, the tandem node makes bandwidth reservations on input and output links, increments the hops field in the APR packet and forwards the request to the next appropriate link (step 365). The tandem node sets the path-state to 'path adding' and waits for response from the next link (step 370).

Receiving an Add Path Request packet at a Destination Node

Fig. 4 is a flow chart illustrating the actions performed by a destination node when the destination node receives an APR packet. Initially, the destination node receives the APR packet (step 405). The destination node determines if the request is invalid and contains errors (step 410). If the request is valid, the destination node proceeds with resource checking (step 445). If the request contains errors, the destination node determines if one or more of the parameter fields in the APR packet are valid (step 415). These parameters may include bandwidth, type, CoS, and Link IDs and the like. If one or more of the parameters are invalid, the destination node responds with a NAK (WRONG PARAMETERS) upstream (step 420). The destination node determines if the APR packet contains a path length that does not match a hops field of the APR packet (step 425). If the path length does not match the hops field, the destination node sends a NAK (INVALID PATH) (step 430). The destination node determines if the APR packet is received in an invalid path state (i.e., the request arrives while the path is in restoring or deleting states) (step 435). If the APR packet arrives during an invalid path state, the destination node responds with a NAK (WRONG STATE) (step 440).

If no errors are found in the APR packet (step 410), the destination node determines if sufficient resources are available to support the virtual path (step 445).
 20 If sufficient resources (such as memory, bandwidth and like) are not available, the destination node responds with a NAK (NO RESOURCES) (step 450). If sufficient resources are available, the destination node allocates resources and makes appropriate connections in a cross-connect matrix for the virtual path (step 460). The destination node reformats an Add Path Response packet with a list of ports assigned to the virtual path and sends the Add Path Response packet upstream (step 470).

Add Path Response Packet

Fig. 5 illustrates an example of an Add Path Response packet (“response packet”) 500. The command-specific contents of response packet 500 are similar to those of APR packet 200. In response packet 500, a list of Port IDs 510 is added.
 30 Every node that receives response packet 500 adds a list of Port IDs to the packet.

These ports are assigned on the upstream link of the node. The Port IDs are local port IDs for the node that is the next node to receive the response packet 500. The number of ports assigned for the VP is same as the number of bandwidth units requested (e.g., in terms of STS-48 granularity).

- 5 For positive responses, the responding node copies the contents of the APR packet 200 and appends a port index 505 (e.g., Port IDs 510(1)–(n)) and decrements hops field 280. For negative responses, the responding node copies the contents of the APR packet and instead of a list of assigned Port IDs, the node appends a reason code for rejection in place of port index 505 and decrements hops field 280.

10 Receiving Add Path Response packet at the Source Node

- Fig. 6A is a flow chart illustrating the actions performed by a source node when the source node receives a response packet. Initially the node receives the response packet (step 605). The source node determines if the response packet contains any errors (step 610). If the response packet contains any error (i.e., the response is a negative response), the source node determines if responses from all previously generated APR packets have been received (step 615). If the responses from all previously generated APR packets have been received, the source node generates a network alarm (step 620). If the responses from all previously generated APR packets have not been received, the source node ignores the response and takes no action (step 625). The source node waits until the all the responses from previously generated APR packets have been received (step 630).

- If no errors are indicated in the response packet (i.e., the response is a positive response), the source node determines if the port index in the response packet contains valid Port IDs (step 635). If the port index contains one or more invalid port IDs, the source node generates a network alarm (step 640). If the port IDs in the port index are valid, the source node terminates any timer the source node had to monitor the response time (step 645). The source node stops generating new APRs (step 650). The source node ignores any further response from previously generated APR packets (step 655). The source node restores the virtual path by provisioning the virtual path on the ports allocated by the tandem nodes and the destination node (step 660).

Receiving Add Path Response packet at the Tandem Node

Fig. 6B is a flow chart illustrating the actions performed by a tandem node when the tandem node receives a response packet. Initially the tandem node receives the response packet (step 662). The tandem node determines if the response packet contains any errors (step 665). If the response packet contains any error (i.e., the response is a negative response), the tandem node determines if the attempt count field of the response packet is lower than the attempt count field of the last APR packet forwarded by the tandem node (step 670). If the attempt count field of the response packet is higher than the attempt count field of the last APR packet forwarded by the tandem node, the tandem node releases the resources allocated to the VP (step 673). The tandem node forwards a negative response upstream (step 675). If the attempt count field of the response packet is lower than the attempt count field of the last APR packet forwarded by the tandem node, the tandem node drops the response packet and takes no action (step 678). The tandem node drops the packet because the tandem node may, in future, receive a response packet with no errors (i.e., a positive response) to a later-forwarded APR packet. The tandem node waits for next response packet (step 680).

If the response packet contains no errors, the tandem node determines if the port index in the response packet contains valid port IDs (step 683). If the port index in the response packet contains one or more invalid port IDs, the tandem node releases the resources allocated to the VP (step 685). The tandem node forwards a negative response, NAK (INVALID PORT), upstream (step 688). If the port IDs in the response packet are valid, the tandem node terminates any timer the tandem node had to monitor the response time (step 690). The tandem node makes appropriate connections in the cross-connect matrix (step 693). After making the connections, the tandem node forwards the response packet to next node in the VP (step 695).

An Example of Path Restoration Using Dynamic Unicast Method

The following description is intended to be illustrative of the invention and should not be taken to be limiting. Other embodiments within the scope of the present invention are possible.

Fig. 7 illustrates an example network 700 in which a VP fails and then is restored. Network 700 comprises seven nodes, nodes 701-707. Each node is connected to an adjacent node by an optical link. Each optical link includes multiple physical ports. The failed VP (VP 0) is restored by provisioning VP 0 over nodes 5 701, 702, 705 and 707 using links 712, 725, and 757. Node 701 is the source (or origin) node of VP 0 and node 707 is the destination (or target node) of VP 0. Node 701 (the source node) generates and sends an APR packet to restore VP 0 starting at node 701 and ending at node 707.

Add Path Request Packet Flow

10 In the present example, the configuration requires that the path through nodes 702 and 705 be considered to restore VP 0. Node 701 (the source node) generates an APR packet that traverses through nodes 702 and 705 to reach node 707 (the destination node). Node 701 sets the path length field to three to indicate the number of hops to be traversed by the APR packet. Node 701 updates the path field by adding link 712 to the list. Link 712 is used to send the APR packet to node 702. Node 701 sends the APR packet to node 702. Node 702 increments the hops field (e.g., by 1) to indicate the number of hops traversed by the APR packet and adds the next link ID, link 725, in the path field and forwards the APR Packet to node 705. Node 705 adds link ID, link 757, in the path field, increments the hops field (e.g., by 1) and forwards 15 the APR packet to node 707, the destination node. Table 1 shows the values of some 20 of the fields in the APR packet at each node.

Packet Flow	Path Length field value	Hops field value	Path field value
Node 701 → 702	3	0	Link 712
Node 702 → 705	3	1	Link 725
Node 705 → 707	3	2	Link 757

Table 1. Some of the field values for the APR packet.

25

Add Path Response packet flow

Node 707 determines from the APR packet header that node 707 is the target (destination) node for the APR packet. Node 707 generates a response packet and copies the contents of APR packet to the response packet and decrements the hops

field. Node 707 allocates one or more ports for VP 0 and appends the remote port ID of each allocated port to the port index field of the response packet. The remote port ID of allocated port at node 707 is the local ID of that port at node 705. For example, node 707 allocates ports with local IDs 6 and 3 to VP 0. However, port 6 at node 707
 5 is recognized as port 1 by node 705 and port 3 at node 707 is recognized as port 5 by node 705. Thus, for local port 6 at node 707, the remote port ID is 1 and for local port 3 at node 707, the remote port ID is 5. By appending remote port IDs of corresponding local ports to the response packet, node 707 indicates which ports should be allocated by node 705. Each node exchanges the information regarding the remote IDs of local ports during a periodic network topology update process. The port list in the response packet is updated at every hop along the path to the source node.
 10 Table 2 shows the values of some of the fields in the response packet at each node.

Packet Flow	Path Length field value	Hops field value	Path field value	Port Index value
Node 707 → 705	3	2	Link 757	1, 5
Node 705 → 702	3	1	Link 725	7, 10
Node 702 → 701	3	0	Link 712	4, 9

15

Table 2. Some of the field values for response packet.

After receiving the response packet, node 701 (the source node) provisions VP 0 to the new physical path including links 712, 725 and 757. The new physical path for VP 0 uses ports 4 and 9 on link 712 at node 701, ports 7 and 10 on link 725 at node
 20 node 702, ports 1 and 5 on link 757 at node 705 and corresponding ports 6 and 3 at node 707. The network topology at each node is updated to reflect the new physical path for VP 0.

While particular embodiments of the present invention have been shown and
 25 described, it will be obvious to those skilled in the art that, based upon the teachings herein, changes and modifications may be made without departing from this invention and its broader aspects and, therefore, the appended claims are to encompass within their scope all such changes and modifications as are within the true spirit and scope of this invention. Furthermore, it is to be understood that the invention is solely
 30 defined by the appended claims.